Relief R-CNN : Utilizing Convolutional Feature Interrelationship for Fast Object Detection Deployment

Guiying Li, Junlong Liu, Chunhui Jiang
Liangpeng Zhang and **Ke Tang**
School of Computer Science and Technology
University of Science and Technology of China
Hefei, Anhui 230027,P.R. China
Email: lgy147@mail.ustc.edu.cn

Zexuan Zhu
College of Computer Science and Software Engineering
Shenzhen University
Shenzhen 518060, China
Email: zhuzx@szu.edu.cn

arXiv:1601.06719v3 [cs.CV] 20 Sep 2016

Abstract

R-CNN style models are the state-of-the-art object detection methods, which consist of region proposal generation and deep CNN classification for objects. The proposal generation phase in this paradigm is usually time consuming and not convenient to be deployed on hardwares with limited computational ability. This article shows that the high-level patterns of feature value in deep convolutional feature maps contain plenty of useful spatial information, and propose a simple manually designed approach that can extract the information for fast region proposal generation. The proposed method, dubbed Relief R-CNN(R^2-CNN), is a new deep learning approach for object detection. By extracting positions of objects from high-level convolutional patterns, R^2-CNN generates region proposals and performs deep classification simultaneously based on the same forwarding CNN features, unifies the formerly separated object detection process and speeds up the whole pipeline without additional training. Empirical studies show that R^2-CNN could achieve the best trade-off between speed and detection performance among all the comparison algorithms.

Introduction

One type of the state-of-the-art deep learning methods for object detection is R-CNN (Girshick et al. 2014) and its derivative models (Girshick 2015; Ren et al. 2015). R-CNN consists of two separated procedures: reduced object proposal generation and object classification. The object proposal generation focuses on finding the region of interests(ROIs) (Ren et al. 2015) which may contain objects. The classification phase receives the generated ROIs for a deep CNN (Krizhevsky, Sutskever, and Hinton 2012) to classify these ROIs as specific objects or background.

However, R-CNN may suffer from its low efficiency, not only due to the inelegant combination of ROIs and CNN, but also the high computation cost of ROIs generation and deep CNN process. On the one hand, commonly used deep CNN models are computationally intensive and only efficient on special hardwares(e.g. GPU, TPU) (Krizhevsky, Sutskever, and Hinton 2012; Simonyan and Zisserman 2015; He et al. 2015b; Ioffe and Szegedy 2015; He et al. 2015a; Szegedy et al. 2015). On the other hand, various region proposal methods are also time consuming (Hosang et al. 2015; Hosang, Benenson, and Schiele 2014). Even worse, there are not too many acceleration methods for them. These shortages make a trained R-CNN style model uneasy to be deployed on computationally sensitive platforms.

Many approaches have been proposed to accelerate R-CNN, but they are still not superior enough to be deployed on low-end hardwares with respect to the real-time requirement. Some approaches focus on simplifying the structure of CNN models (Han, Mao, and Dally 2016; Kim et al. 2016). Some focus on adapting the combination of CNN and ROIs, such as Fast R-CNN (Girshick 2015) and Faster R-CNN (Ren et al. 2015). Others focus on reducing the computation of ROI generation, such as some windows based methods(Hosang et al. 2015). Fast R-CNN is an accelerated version of R-CNN. It reconstructs the combination of ROIs and CNN by directly mapping the region proposals to a ROI layer inside the deep CNN model. Faster R-CNN integrates proposal generation process into the Fast R-CNN model by means of a region proposal network(RPN). However, similar to other large scale neural networks, Faster R-CNN is prohibitive for computation sensitive platforms. For ROI generation, compared to traditional grouping based methods (Uijlings et al. 2013; Hosang et al. 2015), some windows based methods are faster, such as Bing (Cheng et al. 2014) and Edgeboxes (Dollár and Zitnick 2015). All these methods can not be efficient deploy on low-end hardwares.

In this paper, we propose Relief R-CNN(R^2-CNN) to speed up the deployment of ROI generation in any trained R-CNN style models without any extra training. R^2-CNN is inspired by the similarity between relief sculptures in real life and feature maps in CNN. Visualization of convolutional layers(Zeiler and Fergus 2014; Simonyan, Vedaldi, and Zisserman 2013; A.Dosovitskiy, J.T.Springenberg, and T.Brox 2015; Mahendran and Vedaldi 2015) has shown that convolutional features with high values in a trained CNN directly map to the recognizable objects on input images. Therefore, R^2-CNN utilizes these convolutional feature interrelationships for region proposal generation. It is done by directly extracting the local region wrapping features with high values as ROIs. This approach is faster than many other methods, since a considerably large part of its computations are comparison operations instead of time consuming multiplication operation(Ren et al. 2015; Jia et al. 2014;

Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

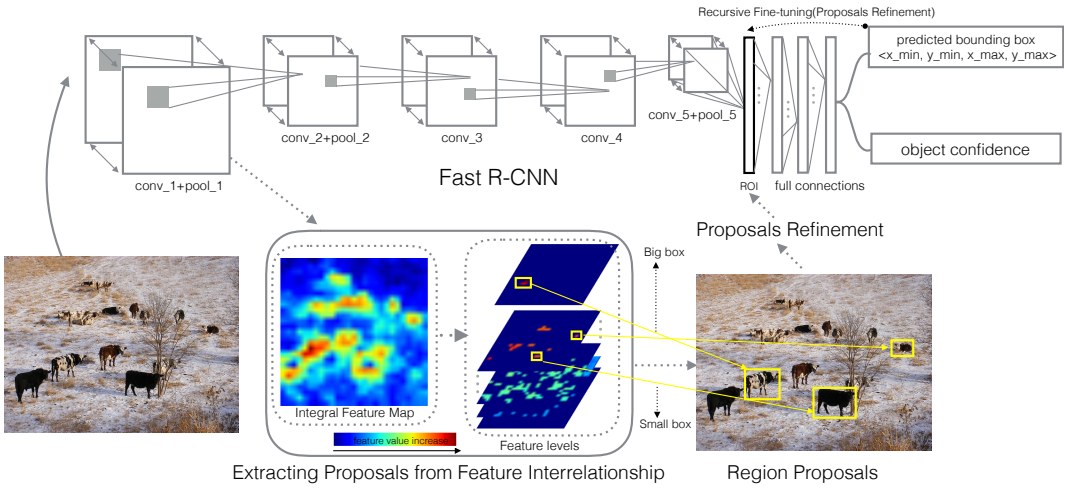

Figure 1: Overview of Relief R-CNN. Generating an **Integral Feature Map** based on feature maps in pool1 layer, followed by separating features of **Integral Feature Map** into different **Feature Levels**. Finally, extracting **Big Boxes** and **Small Boxes** as region proposals and using additional proposal refinement techniques for better performance. The process conducted with solid lines is the procedure of Fast R-CNN, while the process along with dotted lines is the operation of R^2-CNN

Ghodrati et al. 2015; Dai, He, and Sun 2015). Furthermore, R^2-CNN uses the convolutional features produced by CNN for ROI generation, while most of the methods need additional feature extraction from raw images for ROIs. In a word, R^2-CNN could reduce much more computations in ROI generation phase compared to other methods discussed above.

The basic structure of R^2-CNN is shown in Figure 1. A trained Fast R-CNN is the base model which needs acceleration in deployment. R^2-CNN accelerates the ROI generation roughly by 3 steps.

- Firstly, R^2-CNN combines the feature maps in the first convolutional layer as an $Integral\ Feature\ Map$, for the purpose of denoising and simplifying further processing.
- Secondly, R^2-CNN extracts the local regions which contain the features significantly higher than neighbor features as ROIs.
- Thirdly, the generated ROIs are transferred to the ROI layer in Fast R-CNN as inputs.

The main contributions of R^2-CNN can be summarized into 2 points. The first is proposing a real-time region proposal generator with competitive qualities. The second is revealing that, not only the convolutional features values but also the convolutional feature interrelationships in CNN, contain pretty much useful information.

Related Work

Many studies have been devoted to the trade-off between speed and accuracy of region proposal generation(Hosang et al. 2015; Rahtu, Kannala, and Blaschko 2011). Most of them require prior knowledge to formulate rules for region proposals extraction. Two general approaches in those early studies are Grouping based methods and Windows scoring methods (Hosang et al. 2015; Hosang, Benenson, and Schiele 2014).

Grouping based methods, which are based on grouping pixel details for region proposals, are usually time consuming in practice. Selective Search(SS) (Uijlings et al. 2013; Van de Sande et al. 2011) is the state-of-the-art grouping based method, which is highly used in various detection models as baseline (Girshick et al. 2014; Girshick 2015; Hosang et al. 2015). SS utilizes the super-pixels for grouping by manually designed merging process. It merges pixel-details into variant object proposals.

Windows scoring methods, which may need few data-dependent training before testing, are usually faster than grouping based methods, since they can take advantage of some data-dependent priors for locating objects. These methods obtain some coarse window candidates for objects at first. Those coarse windows can be predefined(e.g. Bing) or quickly generated(e.g. EdgeBoxes). Then, these methods use the high level features from images such as objectness (Alexe, Deselaers, and Ferrari 2010) to judge whether an image window contains an object. Objectness (Alexe, Deselaers, and Ferrari 2012; 2010), a well known method, generates an initial window set from the salient locations in an image, then exploits image cues to score the initial window set. Bing (Cheng et al. 2014) presents a very fast approach by only generating proposals with predefined window size, and measures objectness with a simple linear model trained over edges features. Edgeboxes (Zitnick and Dollár 2014; Dollár and Zitnick 2015; 2013) generates limited coarse windows, and then evaluating the windows by object boundary estimation and some fine-tuning techniques.

Grouping based methods and Windows scoring methods do not cause too much time budget in training. Faster R-CNN (Ren et al. 2015), on the contrary, spends lots of time in training to produce a data-dependent region proposal network(RPN) for quick and accurate testing. This technique adopted by RPN is not practical for fast deployment on computation limited equipments, since the speed of neural

network processing on computation limited platforms is restricted. Faster R-CNN provides a more compact combination of ROI generation and object classification, and casts object detection acceleration problem as deep neural network acceleration problem which is still not solved for resource limited environments.

The proposed R^2-CNN in this article, for the purpose of accelerating testing phase without additional training effort, utilizes the network-dependent prior to locate objects in testing. It achieves a better trade-off between time and speed than previous approaches. Further more, it makes proposal generation no more obstacle for platforms with limited computational ability.

Relief R-CNN

In this section we present the details of R^2-CNN. Figure 1 shows the brief structure of R^2-CNN.

General Idea

Observations on convolutional feature maps (Zeiler and Fergus 2014; Mahendran and Vedaldi 2015; Gatys, Ecker, and Bethge 2015) have shown that, the features related to recognizable objects have significant higher values than the nearby features from the context. Similar pattern can be seen in relief: the background is filtered away. The key parts of the vision are portrayed on the canvas. The objects are distinguished by the boundaries which present significant height than nearby elements. The elements in the same key part have similar heights. These similarities inspire us that the relative location information of feature values in feature maps can be treated as a kind of edge information, so the objects could be captured from the convolutional feature maps, rather than the raw image. Therefore, only once feature extraction from image is necessary compared to most of the other methods need twice for CNN process and ROI generation separately. This will dramatically reduce feature extraction time.

In this section, a manually designed proposal extraction method inspired by relief sculpture is proposed. The specific method adopted in this article highly rely on the hyper parameters related to Alexnet (Krizhevsky, Sutskever, and Hinton 2012) used in experiments, but the main idea can be applied to any CNN architectures which fit the observation discussed above.

The main idea of the proposed method is simple. By searching the regions that are significant more salient than nearby context features in convolutional feature maps, we can obtain some edge details of objects, and locate the objects in the source image by utilizing these details. R^2-CNN can be summarized into 5 steps as follow, step 1$\sim$3 generate an initial proposal set, step 4$\sim$5 refine the proposal set:

1 Generating an $Integral\ Feature\ Map$ from the first pooling layer for denoising and fast processing in following steps.

2 Separating the features in the $Integral\ Feature\ Map$ into different $Feature\ Levels$.

3 Grouping all the features in a $Feature\ Level$ as different ROIs.

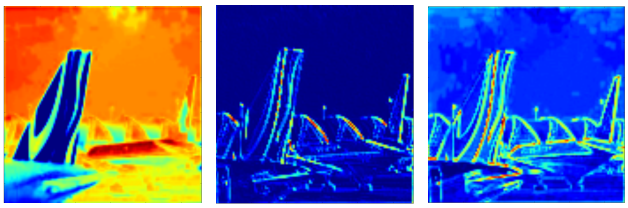

Figure 2: Different appearances of feature maps. Not all the feature maps focus on the objects of interest. *Left* is more interested in context information; *Right* does not filter all the background away; *Middle* is exactly what we want. Although different feature maps have different interests, the boundary features in all these maps are described. The boundary features may not be obvious in global, but significant in local across different maps.

4 Augmenting ROI numbers by *Local Search*.

5 Promoting the performance by *Recursive Fine-tuning*.

Figure 3 shows a simple illustration of steps 1$\sim$3. The following sections will describe these 5 steps in detail.

Initial Proposal Set

Step 1. Integral Feature Map Generation

A synthetic feature map called $Integral\ Feature\ Map$ is generated by adding all feature maps up to one map. It would be a very time consuming task to process all maps in the chosen layer. Furthermore, not all the feature maps contain the information about objects. Some maps are more related to context information(see Fig. 2), so it is meaningless to extract region proposals from such noisy feature maps. Therefore, the $Integral\ Feature\ Map$ could be used for the purpose of eliminating noisy maps and reducing time.

The generation of $Integral\ Feature\ Map$ consists of two steps. Firstly, each feature map is normalized by dividing by its maximal feature value. Secondly, an $Integral\ Feature\ Map$ is generated by adding all the normalized feature maps together.

Step 2. Separating Feature Levels by Feature Interrelationship

Features in the $Integral\ Feature\ Map$ are grouped into several feature levels by uniformly separating their value range. As discussed in **General Idea**, edge details of objects can be described by the relative relationship among features. In this context, the relative relationship between two features indicates how large the gap between values of the two features is. These significant gaps with high magnitudes in feature maps are some kinds of edge details which characterize the contours of objects.

However,there is not a precise threshold to determine whether gaps are "significantly" higher than others. R^2-CNN deals with this problem by separating features into several feature levels, so that the feature values in each feature level are significant higher or lower than those in other levels. It means that features in the same level are highly possible to be part of the same object or similar object. The features in the same feature level form a hierarchical sub feature map, as shown in Figure 3.

These feature levels in an $Integral\ Feature\ Map$ $f_{integral}$ are generated by dividing the value range of all the features into several subranges. Each subrange is a specific level which contains part of features in $f_{integral}$. The number of subrange is a hyper parameter l, R^2-CNN uniformly divides the $f_{integral}$ into l feature levels, see Algorithm 1.

Algorithm 1 Feature Level Separation

Input: Integral Feature Map $f_{integral}$
Input: Feature Level Number l
$\triangleright$ Getting the value range of feature values in $f_{integral}$
Finding the maximal value $value_{max}$ in $f_{integral}$
Finding the minimal value $value_{min}$ in $f_{integral}$
$\triangleright$ Value range is $(value_{max} - value_{min})$
$\triangleright$ Uniformly dividing the value range into l subranges
$stride = (value_{max} - value_{min})/l$
$\triangleright$ features in each subrange form a feature level
for $i = 1 \rightarrow l$ **do**
    Finding features bigger than $value_{min} + (i-1) * stride$ and smaller than $value_{min} + i * stride$ in $f_{integral}$ as $feautre_{level_i}$
    $feature_{level_i}$ is the feature level i for $f_{integral}$
end for
return $< feature_{level_1}, ..., feature_{level_l} >$

Step 3. Bounding Boxes Generation

Bounding boxes for objects are generated by combining the features in each feature level. The techniques for combination are elaborating designed on the observations discussed above.

Figure 3 shows some samples of feature levels generated from the first pooling layer of CaffeNet model(CaffeNet is a caffe implementation of AlexNet (Krizhevsky, Sutskever, and Hinton 2012)). In these samples, each bright pixel in feature levels is a feature. Further more, the position connected features, which are surely belong to the same object, may be an important part of an object or even represent a small object. Therefore, proposals can be generated by assembling these position connected features, or directly using the areas contained such features as proposals. Considering these two approaches of proposal generation, two types of proposals are generated for a specific feature level $feature_{level_i}$, the process is shown in Figure 3:

- Big Box Proposal: This comes from the idea of assembling areas of position connected features. Instead of various possibilities that combine areas into boxes, this approach combines all the areas of position linked features into one **Big Box** for the $feature_{level_i}$. Such a method can deal with the big object which cover a large proportion of target image.
- Small Box Proposals: This comes from the idea of directly using the areas of position connected features as proposals. Firstly, it searches the feature clusters(namely the position connected features) in the given $feature_{level_i}$, and then mapping the feature clusters as **Small Boxes**.

All the **Big Boxes** and **Small Boxes** constitute the initial proposal set.

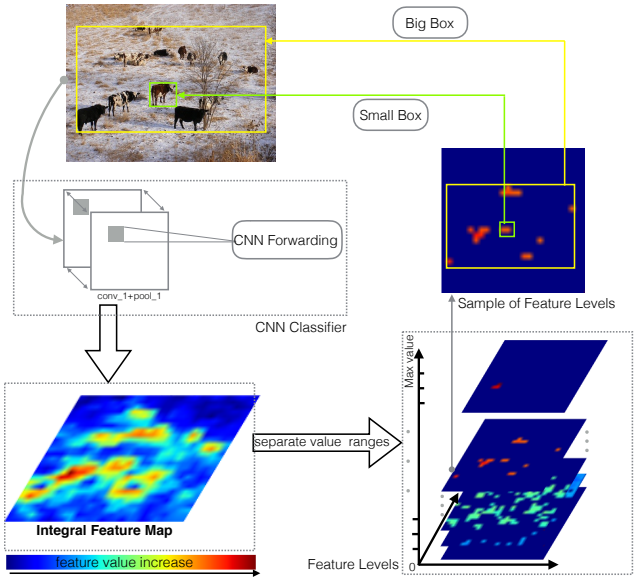

Figure 3: Feature levels separating and Big/Small boxes Extracting. Feature levels separating shows the value distribution of features in a feature map, the same important parts of a feature map are grouped in the same hierarchical sub feature map. One **Big Box** for each sub feature map can be taken as object composed of key parts. There are also many **Small Boxes** in each sub feature map since small objects may appear as feature clusters in the sub map. Feature levels are uniformly separated into ten levels in experiment.

Proposal Refinement

R^2-CNN provides a fast ROI generation for testing while the pre-trained deep model is converged in training by other proposal method(e.g. SS). It is obvious that the accuracy of testing is restricted because of the different proposals distribution between training and testing. Owing to this fact, some fine-tuning methods should be applied to the initial proposal set for a better detection rate.

The purpose of R^2-CNN is to accelerate testing phase without more resource consuming to the whole R-CNN style model. Therefore, a fine-tuning method deployed in testing phase is more reasonable than a deep model retrained by the window candidates from the proposed method. The following section introduces two proposal refinement techniques for the initial proposal set.

Step 4. Local Search

Convolutional features from source image are not produced by seamless sampling. As a result, bounding boxes extracted in convolutional feature maps are quite coarse after mapping to source images. Local Search in width and height is applied to tackle this problem. For each region proposal, local search algorithm needs two scale ratios α and β to generate 4 more proposals by scaling its width and height according to scale ratios. In experiments, α is fixed to 0.8 and β is fixed to 1.5. The Local Search can give about 1.8% mAP improvement in detection performance.

Step 5. Recursive Fine-tuning

Table 1: Testing Time comparison. The object detection model used here is Fast R-CNN. The R^2-CNN needs recursive fine-tuning which makes classification be time-consuming. "Total Time" is the sum of values in "Proposal Time" and "Classification Time" . "*" indicates the runtime reported in (Hosang et al. 2015). Number of recursive loops to R^2-CNN was set to 3. "RPN" is the proposal generation model used in Faster R-CNN. **Bold** means the best time cost.

Methods	Proposal Time(sec.)	Proposals	Classification Time(sec.)	Total Time(sec.)
R^2-CNN	**0.00048**	760.19	0.146	**0.14648**
Bing	0.2*	2000	0.115	0.315
EdgeBoxes	0.3*	2000	0.115	0.415
RPN	1.445	2000	0.115	1.560
Objectness	3*	2000	0.115	3.115
Selective Search	10*	2000	0.115	10.115

A commonly used bounding box refinement technique for R-CNN style models is the bounding box regressor. The trained regressor predicts the corrections for the input proposals so that it will fit the objects much better. The process of a trained regressor in deep model can be formalized as an optimization process $ROI_{out} = Regressor(ROI_{in})$. In this equation, ROI_{in} is the input proposals, and ROI_{out} is the refined proposals which have a better overlapping rate over ground truth bounding boxes.

Fast R-CNN directly uses the ROI_{out} as the predicted object bounding box. This is fine when proposals in training and testing are generated by the same method. However, it is not a stable optimization process from the control system point of view, because it does not provide any feedbacks to verify whether the result is optimal or not, so there is no guarantee of the performance in testing if the method for proposals generation are different from training. Therefore, we developed a close-loop bounding box regressor, namely the *Recursive Fine-tuning* step. This step links the ROI_{out} as the input of $Regressor(ROI_{in})$ again, and only stops when detection performance of ROI_{out} converges.

The recursive fine-tuning is a very simple step. It does not need any changes to existing R-CNN style models, but just a recurrent link from the output of a trained box regressor back to its input. Briefly speaking, it is a trained box regressor wrapped up into a closed-loop system from a R-CNN style model.

Experiments

Setup

In this section, we compared our R^2-CNN with some state-of-the-art object detection methods. All experiments were tested on PASCAL VOC 2007 (Everingham et al. 2014) based on Fast R-CNN model.

For the R^2-CNN in the experiments, the number of recursive loops was set to 3, and the number of feature levels was 10. For easy implementation, we used an individual CaffeNet to do proposals extraction, not exactly as in Fig 1 where the CaffeNet for feature extraction was combined inside the Fast R-CNN.

The proposals of Bing, Objectness, EdgeBoxes and Selective search were the pre-generated proposals published by (Hosang et al. 2015), since the the algorithm settings were the same. The evaluation code used for generating figure 4,5 was also published by (Hosang et al. 2015).

In experiments, the Fast R-CNN (Girshick 2015) was based on CaffeNet which was a Caffe (Jia et al. 2014) version of AlexNet (Krizhevsky, Sutskever, and Hinton 2012). The Fast R-CNN model was trained with SS just the same as in (Girshick 2015). The Faster R-CNN (Ren et al. 2015) used in experiments was based on project py-faster-rcnn (rbgirshick 2016), which was a python version of Faster R-CNN.

Despite the difficulty of Faster R-CNN for low power devices, RPN of Faster R-CNN is still one of the state-of-the-art proposal methods. Therefore, RPN was still adopted in experiments using the same Fast R-CNN model consistent with other methods for detection. The RPN in experiments was trained on the first stage of Faster R-CNN training phases. This paradigm is the unshared Faster R-CNN model mentioned in (Ren et al. 2015). In (Ren et al. 2015), the region proposals need NMS(non-maximum suppression) after generation, which is not a part of RPN model itself and requires more generation time, hence the proposals used here are the raw proposals without NMS.

It should be noticed that all the proposal generation parts were processed in CPU(including R^2-CNN and RPN) while the deep neural networks of classifications were processed in GPU. All the deep neural networks had run on one NVIDIA Titan X GPU card, and the CPU used in the experiments was Intel E5-2650V2 with 8 cores, 2.6Ghz.

Speed Performance

Table 1 contains the results of comparison about time in testing. The testing time is separated into proposal time and classification time. The proposal time is the time cost for proposal generation, and the classification time is the time cost for verifying all the proposals.

It should be noticed that the time costs of convolutional layers are included in classification time, so that the proposal generation time for R^2-CNN in table 1 is only the proposal calculation time based on the convolutional features as inputs. This is because of the shared feature nature of proposal generation and classification in R^2-CNN.

Detection Performance

Table 2 has shown the detection performances of R^2-CNN and other comparison methods. Precision (Özdemir et al.

2010) is a well known metric to evaluate the precision of predictions, mAP(mean Average Precision) is a highly accepted evaluation in the object detection task (Russakovsky et al. 2014). All the results in table 2 were evaluated by the Fast R-CNN detection model based on CaffeNet.

Table 2: Detection performance of R^2-CNN Compare to Other Methods. R^2-CNN shows comparable detection performance. **Bold** results are the results of R^2-CNN

Methods	mAP	mean Precision
Bing	41.2	2
Objectness	44.4	1.7
RPN	50.9	8.2
R^2-CNN	**53.8**	**9.2**
EdgeBoxes	55.5	4.2
Selective Search	57.0	5.9

The empirical results in table 1 and 2 reveal that R^2-CNN outperforms other methods considering the trade-off between time and detection performance. R^2-CNN could achieves a very competitive detection performance compared to state-of-the-art SS and EdgeBoxes with a much more fast speed.

Proposal quality

To evaluate the quality of proposals, two commonly used evaluation metrics (Hosang et al. 2015) are adopted.

1 One is the curves of Recall-to-Proposals under different IoU thresholds. Recall (Özdemir et al. 2010) is a well known evaluation to analyze how many ground truth objects are found.

2 The other evaluation metric is Recall-to-IoU curve.

The metric IoU means intersection over union (Russakovsky et al. 2014), it is an evaluation criterion to measure how similar two boxes are.

Figure 4 contains Recall-to-Proposals curves. It could be found that R^2-CNN was more stable under different IoU thresholds compared to other windows based methods. Bing and Objectness only performed well under the IoU threshold 0.50. EdgeBoxes, in IoU threshold 0.70, also performed a different trend compared to threshold 0.5 and 0.8. R^2-CNN performed more stable than these windows based methods, the differences among R^2-CNN and other methods were continually decreasing as IoU threshold became larger. For SS, it performed as stable as R^2-CNN, but R^2-CNN always got better recall when proposals number was less than about 200 proposals in all the IoU thresholds we evaluated. In summary, on the average recall curve we could find that, R^2-CNN performed better than all other methods.

Figure 5 is Recall-to-IoU curves, it could be found that R^2-CNN had nearly dominated other methods in IoU threshold between 0.5$\sim$0.9, and became the secondary best in IoU threshold 0.9$\sim$1.0.

It should be noticed that R^2-CNN can not control the number of proposals, but it gets the best result with hundreds of proposals while others need thousands. The experiments

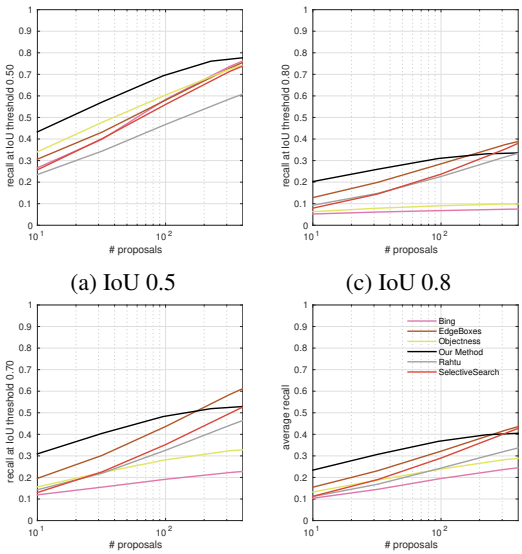

(a) IoU 0.5 (c) IoU 0.8

(b) IoU 0.7 (d) Average Recall

Figure 4: Recall to Proposals curve on VOC07

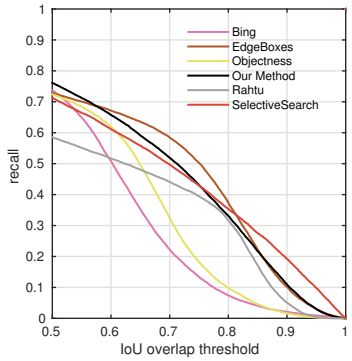

Figure 5: Recall to IoU threshold with 200 proposals in count. R^2-CNN had nearly dominated other methods.

in this section have shown that R^2-CNN can get a very good performance in limit proposals situation with a high speed.

Conclusion

This paper presents a unified object detection model called Relief R-CNN(R^2-CNN), which is based on the similarity between convolutional feature maps and sculpted reliefs. By directly extracting region proposals from convolutional feature interrelationship, namely the location information of salient features in local region, R^2-CNN reduces the ROI generation time in deployment of R-CNN style models. Empirical studies demonstrated that R^2-CNN was faster than previous works with competitive detection performance. Moreover, no additional training budget to original Fast R-CNN baseline model was needed. The results of experiments also revealed a new thinking that information is not only presented as features but also the organization of feature positions in a deep CNN.

References

A.Dosovitskiy; J.T.Springenberg; and T.Brox. 2015. Learning to generate chairs with convolutional neural networks. In *CVPR*. IEEE.

Alexe, B.; Deselaers, T.; and Ferrari, V. 2010. What is an object? In *CVPR*, 73–80. IEEE.

Alexe, B.; Deselaers, T.; and Ferrari, V. 2012. Measuring the objectness of image windows. *PAMI* 34(11):2189–2202.

Cheng, M.-M.; Zhang, Z.; Lin, W.-Y.; and Torr, P. 2014. Bing: Binarized normed gradients for objectness estimation at 300fps. In *CVPR*, 3286–3293. IEEE.

Dai, J.; He, K.; and Sun, J. 2015. Convolutional feature masking for joint object and stuff segmentation. *CVPR*.

Dollár, P., and Zitnick, C. L. 2013. Structured forests for fast edge detection. In *Proceedings of the IEEE International Conference on Computer Vision*, 1841–1848.

Dollár, P., and Zitnick, C. L. 2015. Fast edge detection using structured forests. *IEEE transactions on pattern analysis and machine intelligence* 37(8):1558–1570.

Everingham, M.; Eslami, S. A.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2014. The pascal visual object classes challenge: A retrospective. *IJCV* 111(1):98–136.

Gatys, L. A.; Ecker, A. S.; and Bethge, M. 2015. A neural algorithm of artistic style. *arXiv preprint arXiv:1508.06576*.

Ghodrati, A.; Diba, A.; Pedersoli, M.; Tuytelaars, T.; and Van Gool, L. 2015. Deepproposal: Hunting objects by cascading deep convolutional layers. In *ICCV*, 2578–2586.

Girshick, R.; Donahue, J.; Darrell, T.; and Malik, J. 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 580–587. IEEE.

Girshick, R. 2015. Fast r-cnn. In *ICCV*.

Han, S.; Mao, H.; and Dally, W. J. 2016. Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding. *ICLR* 2.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015a. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015b. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification.

Hosang, J.; Benenson, R.; and Schiele, B. 2014. How good are detection proposals, really? In *BMVC*.

Hosang, J.; Benenson, R.; Dollár, P.; and Schiele, B. 2015. What makes for effective detection proposals? *PAMI*.

Ioffe, S., and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In Blei, D., and Bach, F., eds., *ICML*, 448–456. JMLR Workshop and Conference Proceedings.

Jia, Y.; Shelhamer, E.; Donahue, J.; Karayev, S.; Long, J.; Girshick, R.; Guadarrama, S.; and Darrell, T. 2014. Caffe: Convolutional architecture for fast feature embedding. In *ACM Multimedia*, 675–678. ACM.

Kim, Y.-D.; Park, E.; Yoo, S.; Choi, T.; Yang, L.; and Shin, D. 2016. Compression of deep convolutional neural networks for fast and low power mobile applications. *ICLR*.

Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *NIPS*, 1097–1105.

Mahendran, A., and Vedaldi, A. 2015. Understanding deep image representations by inverting them. *CVPR*.

Özdemir, B.; Aksoy, S.; Eckert, S.; Pesaresi, M.; and Ehrlich, D. 2010. Performance measures for object detection evaluation. *Pattern Recognition Letters* 31(10):1128–1137.

Rahtu, E.; Kannala, J.; and Blaschko, M. 2011. Learning a category independent object detection cascade. In *ICCV*, 1052–1059. IEEE.

rbgirshick. 2016. py-faster-rcnn project. `https://github.com/rbgirshick/py-faster-rcnn`.

Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 91–99.

Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2014. Imagenet large scale visual recognition challenge. *IJCV* 1–42.

Sermanet, P.; Eigen, D.; Zhang, X.; Mathieu, M.; Fergus, R.; and LeCun, Y. 2013. Overfeat: Integrated recognition, localization and detection using convolutional networks. *arXiv preprint arXiv:1312.6229*.

Simonyan, K., and Zisserman, A. 2015. Very deep convolutional networks for large-scale image recognition. *ICLR*.

Simonyan, K.; Vedaldi, A.; and Zisserman, A. 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. *ICLR Workshop 2014*.

Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going deeper with convolutions.

Uijlings, J. R.; van de Sande, K. E.; Gevers, T.; and Smeulders, A. W. 2013. Selective search for object recognition. *IJCV* 104(2):154–171.

Van de Sande, K. E.; Uijlings, J. R.; Gevers, T.; and Smeulders, A. W. 2011. Segmentation as selective search for object recognition. In *ICCV*, 1879–1886. IEEE.

Zeiler, M. D., and Fergus, R. 2014. Visualizing and understanding convolutional networks. In *ECCV*. Springer. 818–833.

Zitnick, C. L., and Dollár, P. 2014. Edge boxes: Locating object proposals from edges. In *ECCV*. Springer. 391–405.